C. Highlights of Research Progress

C.1 SECS Project Developments

The focus of our work this year on SUMEX was conversion of our research programs from the SUMEX DEC-20 to a VAX 11/750 located in our research group. We were restricted to a 3% maximum cpu utilization on SUMEX which effectively precluded significant production work on SUMEX. We completed moving all files from SUMEX to our VAX 11/750 31 March 1985.

C.1.a SECS on VAX

The majority of the SECS program has been converted to the VAX in Fortran and is operational. A graphic driver for the Evans & Sutherland PS330 display system has been also added. New chemical transforms in heterocyclic chemistry have been written and debugged. Through our collaborations in Japan 1000 new chemical transforms have been added using an automatic ALCHEM transform writing program. All new developments in the SECS program will occur on the VAX version.

C.2 XENO Program Developments

The metabolic fate of various compounds in the human body is extremely complex, yet extremely important for it is known that through metabolism certain otherwise harmless compounds are converted into toxic and possibly carcinogenic agents. Because of this complexity it is difficult, looking at a given compound, to forecast potential biological activity of that given compound. The objective of this proposal is to develop a practical computer program by which a biochemist or metabolism expert can explore the metabolites of a given compound and be alerted to the plausible biological activity of each metabolite.

C.2.a Evaluation Study

We participated in an evaluation of XENO predictions of metabolism on four pre-manufacturing notice compounds from the U.S. Environmental Protection Agency, Office of Toxic Substances, in comparison with two panels of metabolism experts. These four compounds were selected from a list of six compounds considered by the EPA to be representative of the types and diversity of compounds they must evaluate. The limit of 3% cpu maximum utilization precluded evaluating more compounds.

The predictions of XENO were submitted to a third party as were the predictions from the two other panels of experts. The results from all three groups were then distributed and discussed at a meeting in Washington, DC.

In processing these four examples, the XENO program performed without crashing or errors. The graphical display equipment broke down during the study, but because XENO also permits teletype graphics, we were still able to complete the study. The total computer time used was approximately 15 minutes on a DEC 2060 system which is very little time for such analyses.

The results of the evaluation proved very interesting. The knowledge base of XENO was shown to be missing a couple transforms having to do with cleavage of C-S bonds, disulfide formation, and phosphorylation. These transforms have now been added to XENO, which required about 15 minutes, illustrating the simplicity of augmenting the knowledge base. But beyond the couple missing transforms, XENO correctly included all predictions by the experts and further suggested additional metabolites that might be present which the experts had not included.

XENO agreed with the experts more than the experts agreed with each other. The experts tended to approach the problems very narrowly, with just a few selected pathways. XENO tended to include the results of all the experts, approaching the problem more broadly. If the objective is risk assessment, the latter strategy is preferable. XENO also suggested some reasonable pathways, such as azo reduction in aryl-alkyl azo compounds, but the experts, having never seen results from such a compound concluded that because it had not been reported, it didn't occur. Now however an experimental study of azo reductase has been launched to determine what does happen with aryl-alkyl azo compounds.

Finally, as might be expected, the experts were biased against the computer expert system, and had greater difficulty seeing its potential than others involved in the risk assessment process.

C.2.b Molecular Model Builder

Over the past year we have begun a new project, to replace the molecular model builder in XENO with a faster and more general one. This will allow steric evaluation to be done more quickly and accurately. The goal of our project is to build a knowledge based program which can quickly and accurately create three-dimensional molecular models of organic molecules. Unlike other numerically oriented modelling programs, our program utilizes a large body of existing conformational data to infer preferred geometries. This knowledge base is the Cambridge Crystal file, which contains x-ray determined geometries for over 20,000 organic compounds.

The design work for the program was completed during the past year and now we are at the early stages of implementation. The program consists of the following individual modules:

1. A graphical front end facilitates input into the program and the display of results. The graphics package is a flexible visual tool for the chemist and runs on an Evans and Sutherland PS300 linked asynchronously to our VAX 750. It allows the easy construction and manipulation of both two-dimensional and crude three-dimension structures.

2. A perception module perceives the input structure for atom types, bond types, stereochemistry, bonding configuration, rings, and ring assemblies.

3. A search strategy generation module uses the perception data to formulate hierarchical rules, constraints, and goals used in searching the data base for possible structural knowledge to be used in model construction. Generation of the search strategy can be interactively guided by constraints and priorities defined by the user.

4. A construction module applies the knowledge found using a set of attachment rules and attempts to construct models which meet the initial constraints.

5. An evaluation module evaluates the models generated to determine the confidence level for the three-dimensional accuracy of each part of the model. This evaluation is based on criteria such as degree of analogy between previous precedent and current model.

Currently, the first, second and fourth modules of the program have been implemented and are being tested.

C.2.c Collaborative Efforts.

The co-operation between the groups at the University of Lund and the University of California Santa Cruz continues to prove fruitful for both parties. The SECS program, which was implemented in Sweden by Dr. Robert E. Carter after his visit to Santa Cruz in 1982, is still being used by both graduate and undergraduate students. Currently, SECS is hosted on a PDP-10 which is located 200 miles to the north of Lund. However, Lund is going to lose access to this machine in the foreseeable future. Fortunately, both Lund and Santa Cruz have purchased VAX machines, and Prof. Wipke has indicated that Lund will receive a VAX version of SECS in the near future.

Further cooperation was accomplished this winter when Dr. Dolata, formerly of Santa Cruz, and now at Lund, visited Prof. Wipke. Since Lund had obtained its VAX about 6 months previous to Santa Cruz, they had had time to build a repertory of useful programs and procedures. These were installed on the Santa Cruz VAX, thus improving the programming environment substantially.

In addition, Dr. Dolata gave a seminar on the current work in conformational analysis by symbolic reasoning which is under investigation at Lund, and received many thoughtful and helpful insights. A copy of the WIZARD conformational analysis system was provided for examination by the Santa Cruz group. Additionally, several papers to be published by Wipke and Dolata were discussed, and work was started on these papers.

Finally, with the upcoming installation of UUCP net on Lunds Vax, communication between UCSC and LU should be facilitated, so that even closer cooperation can be achieved.

The SECS project continues to have collaborations with the pharmaceutical industry which is adding chemical transforms and doing some joint program development, for example, Dr. Yanaka continued work started at Santa Cruz after he returned to Kureha Chemical in Japan and a paper has been prepared on that work.

In addition to collaboration with the SECS project, Dr. David Rogers at the University of Michigan writes: The SUMEX-AIM site has been a useful and necessary link for our AI research group at the University of Michigan to the ARPAnet community. Our work is an attempt to build a working system based on emergent structure appearing as the result of the statistical interaction of low-level subcognitive units; our work is being done on a network of SUN microcomputers using Franz Lisp. We appreciate the existence of SUMEX-AIM as an assist at keeping abreast with work at Stanford and other ARPA sites.

D. List of Current Project Publications

1. Wipke, W.T., and Rogers, D.: *Artificial Intelligence in Organic Synthesis. SST: Starting Material Selection Strategies. An Application of Superstructure Search.* J. Chem. Inf. Comput. Sci., 24:1 71-81, 1984.

2. Wipke, W.T., and Rogers, D.: *Rapid Subgraph Search Using Parallelism* J. Chem. Inf. Comput. Sci., 24:4 255-262 (1984).

3. Wipke, W.T.: "An Integrated System for Drug Design" in *The Aster Guide to Computer Applications in the Pharmaceutical Industry* Aster Publishing Co., Springfield, Oregon, 1984, pp 149-166.

4. Wipke, W.T.: *Computer Modeling in Research and Development*, Cosmetics and Toiletries, 99:Oct 73-82 (1984).

5. Wipke, W.T.: *Computer-Assisted Design of Organic Synthesis. ALCHEM: A Language for Representing Chemical Knowledge*, J. Chem. Info. Comput. Sci., 24, 0000 (1985).

6. Johnson, C.K., Thiessen, W.E., Burnett, M.N., Condran, P. Ronlan, A., Yanaka, M. and Wipke, W.T.: *Systematic derivation of chemical procedures for transforming surplus hazardous chemicals to useful products, J. of Hazardous Materials.* (In press, the appearance of this article has been delayed by Oak Ridge.)

7. Dolata, D.P.: *QED: Automated Inference in Planning Organic Synthesis* (Ph.D. dissertation). University of California, Santa Cruz, 1984.

8. Rogers, D.: *Artificial Intelligence in Organic Chemistry. SST: Starting Material Selection Strategies* (Ph. D. dissertation). University of California, Santa Cruz, 1984.

F. Research Environment

At the University of California, Santa Cruz, we have been previously connected to the SUMEX-AIM resource by a 4800 baud multiplexed leased line. Now we have disconnected that line and are using a VAX 11/750 as our host computer running the VMS operating system. We have a PS300 black/white vector graphic display which is driven by a serial line to the VAX. The SECS laboratory is located in 125 Thimann Laboratories, adjacent to the synthetic organic laboratories at Santa Cruz.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

SECS had been available in the GUEST area of SUMEX for casual users. SECS and XENO are no longer available through SUMEX. Access now must be through UCSC or by installation on the user's own computer.

Communication between SECS collaborators is facilitated by using SUMEX message drops, especially when time differences between the U.S. and Europe and Australia makes normal telephone communication difficult.

B. Examples of Cross-fertilization with other SUMEX-AIM Projects

The AILIST bulletin board has been used extensively for interacting with many projects and locating references for further information related to program design and AI technology. There are no longer any other chemical or biochemical projects on SUMEX so our interaction with the community is limited to AI technology interchange.

C. Critique of Resource Services

SUMEX-AIM gives us at UCSC, a small university, the advantages of a larger group of colleagues, and interaction with scientists all over the country. Since 1 April 1984, the computer response time has been very poor for the SECS project because our project was put in a separate class with a 3% cpu limitation. This was a very severe restriction which prevented short usage peaks from being averaged with other users. Projects in their final year should not be so restricted.

D. Collaborations and Medical Use of Programs via Computers other than SUMEX

SECS 2.9 has been installed on the CompuServe computer networks for the past four years so anyone can access it without having to convert code for their machine. This has proved very useful as a method of getting people to experiment with this new technology. SECS also resides on the Medicindat machine at the University of Gothenborg, Sweden, and is available all over Sweden by phone. Similarly in Australia, SECS resides at the University of Western Australia and is available throughout Australia over CSIRONET. SECS has been installed at two locations in Japan. FSECS has been installed on a DEC-10 at Oak Ridge National Laboratory and serves for collaborative development of that approach with Carroll Johnson. PRXBLD has been disseminated to over 60 sites on various types of computers including DEC-10, DEC-20, IBM, VAX, PRIME, FUJITSU and Honeywell.

III. RESEARCH PLANS (4/85-4/86)

A. Near-Term Project Goals and Plans

Our planned use of the SUMEX resource is simply for message communication with collaborators. We will continue developing the SECS and XENO projects on the VAX 11/750 and incorporate graphics with the Evans and Sutherland PS300 system. A proposal is pending to add color displays to this system.

B. Justification and Requirements for Continued Use of SUMEX

We request to have continued access to SUMEX for receiving and sending messages to collaborators and for access to the important bulletin boards maintained on SUMEX. We may also need to retrieve some of our files archived on SUMEX since in moving ten years of research work off SUMEX it is possible we missed some key file which we will not recognize until we need it.

C. Needs Beyond SUMEX-AIM

In addition to our VAX, we are exploring graphic workstations to achieve a distributed environment since the VAX alone loads down very quickly. And we are seeking to add color to our Evans and Sutherland PS300.

D. Recommendations for Community and Resource Development

An important part of medicine is treatment of diseases with drugs. Drugs are chemicals--chemicals that were designed and synthesized by chemists. Since the termination of the DENDRAL project, there seems to be declining support for artificial intelligence applications in chemistry. We feel that support of this area is essential to the advancement of medicine in this country. The lack of chemists on NIH Research Resources computing peer review is contributing to the problem. Application of artificial intelligence in synthesis planning is one of the more successful current applications and it is now a high priority research area in many foreign countries. To maintain our lead in this technology, further funding is required.

Responses to Questions Regarding Resource Future

The SECS group feels that SUMEX should remain a communications center, but there is little need for it to attempt to grow the mainframe in an effort to supply cpu cycles to individual projects. It is now financially feasible for each project to have its own computer. But there is still a need for network access, knowledge sharing, file transfer, etc. SUMEX could serve this networking aspect with considerably less hardware and staff than it now has.

Since SUMEX no longer purports to serve a national community except for communication, there is no justification for continuing to grow the mainframe.

It is hard to see justification for SUMEX to develop workstation software since that is already being done commercially and since a similar proposal to RR to do same from San Diego was disapproved on the basis of it being inappropriate.

We expect to need access to SUMEX for message purpose only. That access is desired for probably two years or more, or until the UC network is operational. Currently much of the UC network is UNIX and VMS people can't currently connect.

Regarding the imposition of fees for service, I think that would be sad. There are already many networks that operate on a fee for service basis, i.e., Source, CompuServe, etc. If SUMEX had to be on a fee for service basis, it is unclear why the service might not better be handled by existing commercial vendors that have customer relations staff. It is unclear also that NIH grants would allow expenses for communication rather than hard computing.

Finally, just a note that I have appreciated the service of SUMEX, and the staff of SUMEX, although I did not appreciate the 3% limit under which we had to work last year.

IV.B.6. SOLVER Project

SOLVER: Problem Solving Expertise

Dr. P. E. Johnson
Center for Research in Human Learning
University of Minnesota

Dr. W. B. Thompson
Department of Computer Science
University of Minnesota

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

This project focuses upon the development of strategies for discovering and documenting the knowledge and skill of expert problem solvers. In the last several years, considerable progress has been made in synthesizing the expertise required for solving extremely complex problems. Computer programs exist with competency comparable to human experts in diverse areas ranging from the analysis of mass spectrograms and nuclear magnetic resonance (Dendral) to the diagnosis of certain infectious diseases (Mycin).

Design of an expert system for a particular task domain usually involves the interaction of two distinct groups of individuals, "knowledge engineers," who are primarily concerned with the specification and implementation of formal problem solving techniques, and "experts" (in the relevant problem area) who provide factual and heuristic information of use for the problem solving task under consideration. Typically the knowledge engineer consults with one or more experts and decides on a particular representational structure and inference strategy. Next, "units" of factual information are specified. That is, properties of the problem domain are decomposed into a set of manageable elements suitable for processing by the inference operations. Once this organization has been established, major efforts are required to refine representations and acquire factual knowledge organized in an appropriate form. Substantial research problems exist in developing more effective representations, improving the inference process, and in finding better means of acquiring information from either experts or the problem area itself.

Programs currently exist for empirical investigation of some of these questions for a particular problem domain (e.g. AGE, UNITS, RLL). These tools allow the investigation of alternate organizations, inference strategies, and rule bases in an efficient manner. What is still lacking, however, is a theoretical framework capable of reducing dependence on the expert's intuition or on near exhaustive testing of possible organizations. Despite their successes, there seems to be a consensus that expert systems could be better than they are. Most expert systems embody only the limited amount of expertise that individuals are able to report in a particular, constrained language (e.g. production rules). If current systems are approximately as good as human experts, given that they represent only a portion of what individual human experts know, then improvement in the "knowledge capturing" process should lead to systems with considerably better performance.

In order to obtain a broad view of the nature of human expertise, the SOLVER project includes studies in a variety of complex problem solving domains in addition to medicine. These include law, auditing, business management, plant pathology, and expert system design. We have observed that despite the apparent dissimilarities in these problem solving areas there is reason to believe that there are underlying principles of expertise which apply broadly. Our project seeks to investigate these principles and to create tools to make use of that knowledge in practical expert systems.

B. Medical Relevance and Collaboration

Much of our research has been and will continue to be directly focused on medical AI problems. GALEN, our experimental expert system in pediatric cardiology, is achieving expert levels of performance. Dr. Connelly is initiating a project to develop an expert system based platelet transfusion therapy monitoring program. Dr. Spackman is completing a doctoral thesis on the automated acquisition of rule knowledge in medical microbiology.

Some of our research has focused on problems in diagnostic reasoning and expertise in domains other than medicine. However, our experience indicates that principles of expertise and relevant knowledge engineering tools can cut across task domains. GALEN is demonstrably a useful expert system implementation tool designed in the medical diagnostic task domain. Developments from our work in other domains affecting problems such as automated knowledge acquisition through rule induction and reasoning by analogy will have medical relevance.

Collaboration with Dr. James Moller in the Department of Pediatrics, Dr. Donald Connelly in the Department of Laboratory Medicine, at the University of Minnesota. Dr. Connelly has become a SUMEX user and is teaching a course in medical informatics. He has also initiated a project to create an expert system in platelet transfusion therapy. Collaboration with Dr. Eugene Rich and Dr. Terry Crowson at St. Paul Ramsey Medical Center. Dr. Kent Spackman is a post-doctoral fellow in medical informatics who is completing a Ph.D. thesis in Artificial Intelligence. Dr. Spackman is a resident at the University of Minnesota Hospitals and collaborates with the SOLVER project.

C. Highlights of Research Progress

Accomplishments of This Past Year -- Prior research at Minnesota on expertise in diagnosis of congenital heart disease has resulted in a theory of diagnosis and an embodiment of that theory in the form of a computer simulation model, *Galen*, which diagnoses cases of congenital heart disease [Thompson, Johnson & Moen, 1983]. Continuing development and research with GALEN have led to results in analyzing Garden Path problems in medical diagnosis. Such problems are ones in which an initial solution is later proved to be incorrect. Successful solution of such problems depends upon rejecting an initial incorrect response in favor of a later appropriate one. Errors in Garden Path Problems are generally not due to a lack of knowledge but rather to a confusion over the conditions under which specific rules apply. GALEN was used to identify and test strategies for avoiding Garden Path errors as well as the specific clinical knowledge needed to overcome Garden Path errors in diagnostic reasoning. [Johnson, Moen, and Thompson, 1985].

Galen is descended from two earlier programs written here at Minnesota: *Diagnoser* and *Deducer* [Swanson, 1977]. Deducer is a program that builds hemodynamic models of the circulatory system that describe specific diseases. The models are built by using knowledge about how idealized parts of the circulatory system are causally related. Diagnoser is a recognition-driven program that performs diagnoses by successively hypothesizing one or more of these models and matching them against patient data. The models that match best are used as the final diagnosis. A series of experiments carried out at Minnesota have shown that Diagnoser/Deducer performs as well (and sometimes better) than expert human cardiologists [Johnson et al., 1981].

Despite their early successes, Diagnoser and Deducer did not have a clear, comprehensible structure that is required for the kind of experiments we wish to perform. Galen was built to remedy this problem, taking advantage of the experience gained in the design of Diagnoser and Deducer. Additional discussion of the structure of GALEN can be found in prior annual reports and in the relevant publications.

To determine the generality of our model of expertise in diagnostic reasoning, we are also investigating domains outside medicine. As with our work in congenital heart disease, we have concentrated on the design of mechanisms for structuring problem specific knowledge and for focusing limited computational resources.

One of the Principal Investigators has published results of a study in Expertise in Trial Advocacy, discussing the significance of current research in expertise in legal problem-solving. [Johnson, Johnson, and Little, 1985] Research on legal expertise in corporate acquisition problems has also been investigated. The results of that research suggest that expert corporate acquisition attorneys differ from novices in their greater reliance on internalized norms, prototypes and heuristics. Both expert and novice attorneys in the study went beyond the information provided in task cues in interpreting and predicting actions and situation scripts in the simulated problems. The subjects reasoned heuristically as well as logically. Differences between attorneys in different specialty areas were not large suggesting that the subjects within a domain of problem solving such as legal reasoning acquire meta level reasoning skills that apply to issues within and outside their areas of specialization.

Research is also being completed in a study of cognitive strategies used in making strategic decisions in business. Corporate acquisitions were again used as the context in which to examine expertise. Twenty-four executive subjects were asked to perform an experimental task in which they evaluate companies as candidates for acquisition. The goals of the research are to test for the existence of specialty-related reasoning strategies and to determine the importance of strategic and financial information in problem formulation, problem structuring and choice of strategies in problem solving.

Research in Progress --

Since human experts are notoriously poor at describing their own knowledge, our work requires the creation of problem solving tasks through which experts can reveal criteria for initiating specific hypotheses and methods for investigating those hypotheses.

Current techniques of representing hypotheses and their expectations for diagnosis do not, however, provide much detailed information about the control processes experts use to guide their reasoning. Such control processes typically incorporate highly refined heuristics about which the experts are almost wholly unaware. New research is being proposed to investigate these control structures in legal reasoning, specifically in reasoning by analogy in appellate decision making. Reasoning by analogy appears to be an important inference tool used by experts in many domains as a fundamental problem solving tool. The ability to form plausible analogies lies at the heart of much of the expert ability to be generative when faced with unfamiliar problems. This research will include the implementation of a cognitive simulation of the reasoning by analogy process based upon data obtained by observation of experts solving problems. The results of the simulation will be validated by comparison with human subject data.

We are also investigating several research questions relevant to the architecture of Galen. We have designed an interface to Galen so that users who are unfamiliar with the inner workings of the program can interactively enter case data. Designing the interface raised questions about what forms of data are necessary to adequately and completely represent all possible cases.

One project to test the extensibility of GALEN into other domains is being conducted by a graduate student in the Graduate School of Management. His thesis, Auditing Internal Controls: A computational model of the review process, includes the construction of a working expert system using GALEN. The objective of this study is to formulate and test a model of the processes employed by audit managers and partners in reviewing and evaluating internal accounting controls.

Another project explores the extension of the GALEN architecture into a problem in plant pathology. The main purpose of this research is to find out how the basic postulates about expert reasoning made in Galen hold in a second diagnostic domain. The problem domain chosen for this purpose is Plant Pathology. In collaboration with Professor Paul Teng of the Plant Pathology Department of the University of Minnesota a prototype knowledge base has been implemented. Currently, the knowledge base can diagnose ten potato diseases and has 124 rules. The system is going through evaluation and fine tuning to bring it up to an expert performance level. This system will be useful in the Extension Service at the Plant Pathology department at the University of Minnesota, which provides diagnostic information to farmers over the phone lines.

Dr. Spackman's thesis is entitled "Induction of classification rules under the guidance of comprehensibility-enhancing logical structures and diagnostic performance goals." The purpose of this research is to study and implement methodologies for the automated generation of comprehensible decision rules from empiric data, with emphasis upon logic-based knowledge representation formats and upon problems drawn from the domain of medicine. This work builds upon some of the machine learning methodologies developed at the University of Illinois by R. S. Michalski and others.

This work addresses two shortcomings of previous work on induction of classification rules. These are, first, lack of comprehensibility of the induced rules, and second, lack of flexibility in specifying the diagnostic performance (sensitivity, specificity, or efficiency) desired for the rules that are to be derived.

Comprehensibility of the derived rules or descriptions can be enhanced by imposing restrictions upon the format which the rules may take. For example, the restriction of rules to a unate boolean function format allows the induction of rules that can often be simplified to a "criteria table" type of representation. The type of diagnostic performance a rule must have will depend upon its purpose, and specifying the purpose may allow inductive inference algorithms to trade off small decrements in diagnostic performance for large increments in comprehensibility, or to increase their robustness in the face of noisy or uncertain data.

Successful development of these techniques will lead to enhanced capabilities for deriving rule bases for expert classification systems from empiric data, and will provide new methods for the conceptual analysis of data.

Preliminary results have been obtained for the problem of deriving rules for the identification of bacteria based upon their biochemical profiles in the medical microbiology lab. Other problem domains under investigation are the analysis and interpretation of endocrine laboratory tests, and the induction of rules for the diagnosis of congenital heart disease, for comparison with the rules used in GALEN.

Research is also under way in methods of automating knowledge acquisition in pediatric cardiology. This is being done as thesis research by Paul Krueger. The objective of the research is to design, implement, and test a computerized procedure to derive from examples a nonmonotonic set of rules for an expert classification system. Systems using such rules are generally more efficient than those using monotonic classification processes and more closely approximate psychological models as well.

The research proposes a process for automated learning of preliminary rulebases subject to a set of efficiency constraints which are consistent with a formally defined, psychologically plausible model of classification. The constraints include an upper bound on the amount of information required to explain observations not accounted for by the current set of beliefs, and a lower bound on the degree of inconsistency allowed in the knowledge base at any given time. It will be shown that these constraints can be used to guide the automated determination of both the content and organization of the rules of expert classification systems. The result is behavior that is more focused and efficient, and more closely duplicates the lines of reasoning of domain experts.

A representational formalism for classification knowledge bases based upon a nonmonotonic logic of belief called "autoepistemic logic" (Moore, 1985) is proposed. Having thus defined a representation for the knowledge base the research will propose a methodology for instantiating its concepts within a given application domain. The general approach is to use heuristics to identify from a set of input examples various contextual situations that occur and the types of rules to associate with them. The rule acquisition module (RAM) is then tested in two different application domains. The resulting expert systems will be evaluated for correctness of classification and similarity of their lines of reasoning with those of human experts.

The major conclusion of the research is that constraints similar to those observed in expert human classification processes can be used to guide the empirical induction of efficient expert system rulebases. Supporting this conclusion is the elucidation of a formal nonmonotonic model of classification, and the design and subsequent testing of the Rule Acquisition Module and expert systems derived by it.

D. List of Relevant Publications

1. Connelly, D. and Johnson, P.E.: *Medical problem solving.* Human Pathology, 11(5):412-419, 1980.

2. Elstein, A., Gorry, A., Johnson, P. and Kassirer, J.: *Proposed Research Efforts.* IN D.C. Connelly, E. Benson and D. Burke (Eds.), CLINICAL DECISION MAKING AND LABORATORY USE. University of Minnesota Press, 1982, pp. 327-334.

3. Feltovich, P.J.: *Knowledge based components of expertise in medical diagnosis.* Learning Research and Development Center Technical Report PDS-2, University of Pittsburgh, September, 1981.

4. Feltovich, P.J., Johnson, P.E., Moller, J.H. and Swanson, D.B.: *The Role and Development of Medical Knowledge in Diagnostic Expertise.* IN W. Clancey and E.H. Shortliffe (Eds.), READINGS IN MEDICAL AI, Addison-Wesley, 1984, pp. 275-319.

5. Johnson, P.E.: *Problem Solving.* IN ENCYCLOPEDIA OF SCIENCE AND TECHNOLOGY, McGraw-Hill (in press).

6. Johnson, P.E., Moen, J.B., and Thompson, W.B.: *Garden Path Errors in Medical Diagnosis.* IN Bloc, L. and Coombs, M.J. (Eds.), COMPUTER EXPERT SYSTEMS, Springer-Verlag (in press).

7. Johnson, P.E.: *Cognitive Models of Medical Problem Solvers.* IN D.C. Connelly, E. Benson, D. Burke (Eds.), CLINICAL DECISION MAKING AND LABORATORY USE. University of Minnesota Press, 1982, pp. 39-51.

8. Johnson, P.E.: *What kind of expert should a system be?* J. Medicine and Philosophy, 8:77-97, 1983.

9. Johnson, P.E., *The Expert Mind: A new Challenge for the Information*

Scientist IN Th. M. A. Bemelmans (Ed.), INFORMATION SYSTEM DEVELOPMENT FOR ORGANIZATIONAL EFFECTIVENESS, Elsevier Science Publishers B. V. (North-Holland), 1984.

10. Johnson, P.E., Severance, D.G. and Feltovich, P.J.: *Design of decision support systems in medicine: Rationale and principles from the analysis of physician expertise.* Proc. Twelfth Hawaii International Conference on System Science, Western Periodicals Co. 3:105-118, 1979.

11. Johnson, P.E., Duran, A., Hassebrock, F., Moller, J., Prietula, M., Feltovich, P. and Swanson, D.: *Expertise and error in diagnostic reasoning.* Cognitive Science 5:235-283, 1981.

12. Johnson, P.E. and Hassebrock, F.: *Validating Computer Simulation Models of Expert Reasoning.* IN R. Trappl (Ed.), CYBERNETICS AND SYSTEMS RESEARCH. North-Holland Publishing Co., 1982.

13. Johnson, P.E. and Thompson, W.B.: *Strolling down the garden path: Detection and recovery from error in expert problem solving.* Proc. Seventh IJCAI, Vancouver, B.C., August, 1981, pp. 214-217.

14. Johnson, P.E., Hassebrock, F. and Moller, J.H.: *Multimethod study of clinical judgement.* Organizational Behavior and Human Performance 30:201-230, 1982.

15. Moller, J.H., Bass, G.M., Jr. and Johnson, P.E.: *New techniques in the construction of patient management problems.* Medical Education 15:150-153, 1981.

16. Sedlmeyer, R.L., Thompson, W.B. and Johnson, P.E.: *Knowledge-based fault localization in debugging.* The Journal of Systems and Software, vol. 3, no. 4 (Dec 83) pp. 301-307, Elsevier.

17. Sedlmeyer, R.L., Thompson, W.B. and Johnson, P.E.: *Diagnostic reasoning in software fault localization.* Proc. Eighth IJCAI, Karlsruhe, West Germany, August, 1983.

18. Smith, K.A., Farm, B., Johnson, P.E.: *Surface: A prototype expert system for selecting surface analysis techniques.* Proceedings of IEEE Conference on Computers and Comm., 1985.

19. Swanson, D.B.: *Computer simulation of expert problem solving in medical diagnosis.* Unpublished Ph.D. dissertation, University of Minnesota, 1978.

20. Swanson, D.B., Feltovich, P.J. and Johnson, P.E.: *Psychological Analysis of Physician Expertise: Implications for The Design of Decision Support Systems.* In D.B. Shires and H. Wold (Eds.), MEDINFO77, North-Holland Publishing Co., Amsterdam, 1977, pp. 161-164.

21. Thompson, W.B., Johnson, P.E. and Moen, J.B.: *Recognition-based diagnostic reasoning.* Proc. Eighth IJCAI, Karlsruhe, West Germany, August, 1983.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

A. Medical Collaborations and Program Dissemination via SUMEX

Work in medical diagnosis is carried out with the cooperation of faculty and students in the University of Minnesota Medical School and St. Paul Ramsey Medical Center.

B. Sharing and Interactions with Other SUMEX-AIM Projects

William Clancey, Stanford University, acted as a reviewer of the MEIS Intelligent Systems Project in September, 1984 at the University of Minnesota. The Principal Investigators in the SOLVER project are also principal investigators in that project.

Paul Johnson was a panel member at the SUMEX-AIM conference in Columbus, Ohio in 1984. Dr. Connelly and two graduate students associated with the SOLVER PROJECT also attended the conference.

III. RESEARCH PLANS

A. Project Goals and Plans

Near term -- Our research objectives in the near term can be divided in three parts. First, we are committed to the design, implementation, and evaluation of Galen, as described above. We have completed an interactive front end so that physicians can directly enter patient data, and Galen's knowledge base is currently being "tuned" with the help of Dr. James Moller, an expert physician collaborator from the University of Minnesota Pediatric Cardiology Clinic, the Diagnoser program, and with expert physicians. We believe that GALEN has passed through phases of expertise assessment and cognitive simulation and that it is now approaching a level of performance that will qualify it as a true expert system. An objective now is to extend the explanation capability of GALEN. We are initiating a new investigation into two aspects of expert problem solving that relate to the interaction between a problem solving system and its environment: *"query generation"* and *explanation*. Some simple expert systems proceed from a fixed set of input data to an evaluation of that data. For most problem domains, however, the space of possibly relevant information is large, and some or all of this information may have costs associated with its acquisition. Thus, computational and other costs can be reduced by some mechanism which intelligently selects appropriate queries designed to solicit information that is relevant and cost effective in terms of the problem being solved. Expert systems for complex problem domains must also be able to generate explanations for their actions. Unless the system operates in an entirely autonomous manner, users must be apprised of the rationale for system actions. There is a particular need for explanations tailored for system users rather than system designers.

Experienced experts are typically quite proficient at asking relevant questions, even when the criteria for relevance is difficult to specify. These experts use heuristics capable of keying on selected aspects of data already examined and on the current problem state in order to select the next needed query. We propose to incorporate these heuristics into a *"query generation knowledge base"* . This knowledge base can be thought of as a form of domain specific meta-knowledge. It contains rules by which the problem state can be efficiently evaluated in order to determine the next course of action. By basing these rules on actual expert knowledge and experience, it will often be possible to bypass the combinatorial complexity associated with either blind search or optimization techniques.

Our approach to explanation starts from the premise that substantially different forms of explanation are required within a single expert system. The type of explanation is distinguished both by the level of sophistication of the person receiving the explanation and by whether that person is principally interested in the specific problem being solved or in the internal working of the expert system. Less sophisticated users of the system are likely to have only a superficial understanding of the nature of the system being diagnosed and will require explanations in terms of simplified system properties with which they are familiar. Expert users will require information about significant details of the state of the system being diagnosed and the causal relationships that connect system state with observable symptoms. Designers and maintainers of the expert system require explanations in terms of the actual lines of reasoning used to arrive at a decision.

We will be focusing principally on providing explanations for system *users* rather than system *designers*. Explanations for users must be phrased in terms of the system being diagnosed. Descriptions of the system itself are more important that descriptions of the reasoning strategies used to understand the system. For example, many diagnostic tasks are efficiently approached utilizing recognition-based reasoning strategies using knowledge arising from empirical association. Experts (or possibly automatic learning systems) learn to associate particular interpretations with particular patterns in the data. For many problem domains, knowledge of this sort is quite powerful, providing accuracy without the complexity associated with causal reasoning. The user of such a system, however, requires explanations in terms of causality. This suggest a two-step process. Problem solving is done using a recognition-based strategy. Explanations are generated by combining the results of this process with additional, causally-based explanation knowledge.

Our second objective consists of making extensions to the knowledge capturing strategies developed in our original work in medical diagnosis. In the near term this work will examine descriptive strategies in which experts attempt to use a formalized language to express what they know (e.g. production rules), observational strategies in which experts perform tasks designed to reveal information from which a theory of task specific expertise can be built, and intuitive strategies in which either experts behave as knowledge engineers or knowledge engineers attempt to perform as pseudo experts. The research projects of Dr. Spackman and Paul Krueger which have been discussed previously are both directed toward this objective.

Our third near term objective will be to investigate one of the central problems of recognition based problem solving, how to classify problems when solving them. Questions related to problem classification which we will be examining include: What patterns do experts and novices detect in a problem that allows them to classify it as an instance of a problem type that is already known? How does an expert make an initial choice of the level of abstraction to be used in solving a problem? How can an expert recover from an initial incorrect choice of levels? How can the difference between causal and prototypic modes of reasoning be modeled as differences in levels of abstraction, and how can a common model for these two types of reasoning be constructed? We will be pursuing these questions in the areas of problem solving like law, auditing, and management, as well as in medicine.

Long range -- Our long range objective is to improve the methodology of the "knowledge capturing" process that occurs in the early stages of the development of expert systems when problem decomposition and solution strategies are being specified. Several related questions of interest include: What are the performance consequences of different approaches, how can these consequences be evaluated, and what tools can assist in making the best choice? How can organizations be determined which not only perform well, but are structured so as to facilitate knowledge acquisition from human experts? In the coming year we will be exploring these questions in areas of design and management as well as in law, management and medicine.

B. Justification and Requirements for Continued SUMEX Use

Our current model development takes advantage of the sophisticated Lisp programming environment on SUMEX. Although much current work with Galen is done using a version running on a local VAX 11/780, we continue to benefit from the interaction with other researchers facilitated by the SUMEX system. We expect to use SUMEX to allow other groups access to the Galen program. We also plan to continue use of the knowledge engineering tools available on SUMEX.

We are working toward a Commonlisp implementation of the GALEN system and expect to rely heavily on Commonlisp for future projects.

One of our students implemented a demonstration legal expert system in EMYCIN using the SUMEX resource, and we still find that the resource is valuable for making available major systems which we do not have locally, such as EMYCIN.

C. Needs and Plans for Other Computing Resources Beyond SUMEX-AIM

Our current grant from MEIS has permitted us to purchase four Perq 2 AI workstations for our Artificial Intelligence laboratory. The availability of Commonlisp on these machines is one reason why we expect to make use of that language in the future.

SUMEX will continue to be used for collaborative activities and for program development requiring tools not available locally.

D. Recommendations for Future Community and Resource Development

As a remote site, we particularly appreciate the communications that the SUMEX facility provides our researchers with other members of the community. We, too, are moving toward a workstation based development environment, but we hope that SUMEX will continue to serve as a focal point for the medical AI community. In addition to communication and sharing of programs, we are interested in development of Commonlisp based knowledge engineering tools. The continued existence of the SUMEX resource is very important to us.

IV.C. Pilot Stanford Projects

Following are descriptions of the informal pilot projects currently using the Stanford portion of the SUMEX-AIM resource, pending funding, full review, and authorization.

In addition to the progress reports presented here, abstracts for each project are submitted on a separate Scientific Subproject Form.

IV.C.1. CAMDA Project

CAMDA Project

CAMDA Research Staff:

Prof. Samuel Holtzman, Co-PI	Engineering-Economic Systems
Prof. Ronald A. Howard, Co-PI	Engineering-Economic Systems
Prof. Ross Shachter	Engineering-Economic Systems
Leonard Bertrand	Engineering-Economic Systems
Jack Breese	Engineering-Economic Systems
Kazuo Ezawa	Engineering-Economic Systems
Keh-Shiou Leu	Engineering-Economic Systems
Seok Hui Ng	Engineering-Economic Systems
Emilio Navarro	Engineering-Economic Systems
Dr. Adam Seiver	Engineering-Economic Systems
Joseph Tatman	Engineering-Economic Systems
Dr. Emmet Lamb	School of Medicine
Dr. Robert Kessler	School of Medicine
Dr. Frank Polansky	School of Medicine

Associated faculty:

Prof. Edison Tse	Engineering-Economic Systems

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The Computer-Aided Medical Decision Analysis (CAMDA) project is an attempt to develop intelligent medical decision systems by combining the descriptive generality of expert-system technology with the normative power of decision analysis.

B. Medical Relevance and Collaboration

The primary effort of the CAMDA project during 1984 and early 1985 has been focused on the design and implementation of RACHEL, an intelligent decision system for infertile couples. This system is designed to help patients and physicians deal with difficult medical treatment choices. RACHEL is being developed in close cooperation with the Engineering-Economic Systems Department, the Obstetrics and Gynecology Department, and the Surgery Department (Urology Division), all at Stanford.

In addition to the development of RACHEL, there are several active research programs within the CAMDA project. One such program is aimed at developing a representation for dynamic decision processes (such as those faced by cancer patients) that do not necessarily satisfy the Markov assumption. Another is concentrating on the development of fast algorithms for the solution of general decision problems.

A recent addition to our research project is a program to design cost-effective strategies for monitoring the recurrence of bladder cancer.

C. Highlights of Research Progress

C.1 Accomplishments this past year

We have successfully implemented a pilot-level version of RACHEL. As we define it, a pilot system is one where the essential algorithms work individually as well as interactively with one another, operating with knowledge that is representative of the system's domain. Such a system lacks two important elements that must exist within a prototype-level implementation: an extensive knowledge base, and a front end usable by trained users who may not be familiar with the details of the system.

As part of the development of RACHEL, we have developed a facility to construct individualized models of the patient's preferences over the set of possible outcomes of an infertility therapy. This facility operates in two consecutive stages. The first stage constructs a parametric model from a library of plausible model elements. A typical consideration at this stage is whether to explicitly account for the patient's lifetime. For instance, a treatment strategy which involves surgery would warrant such explicit consideration, whereas a therapy consisting strictly of drugs would not. The second stage in the preference model development process involves the assessment of specific parametric values. These values are obtained directly from the patient to ensure that the overall preference model genuinely reflects his or her desires.

It is important to note that since the preference model is built to fit the specific needs of each case, the interaction between the patient and the system is short and well-focused. In particular, the patient is only asked to respond to a few (about five to ten) questions. These questions are selected so that their relevance to the case is intuitively obvious from the patient's point of view.

Also as a part of RACHEL, we have developed a knowledge base dealing with the decisions faced by the subset of infertile couples whose inability to conceive has been traced to a blockage of the Fallopian tubes of the female partner. In particular, the knowledge in RACHEL deals with the choice between two important procedures pertinent to this condition: laparotomy and in-vitro fertilization.

Another accomplishment during this past research year has been the improvement of our influence-diagram solution procedure. In its original form, this procedure essentially took a brute-force approach to the solution of well-formed influence diagrams. Although its solutions were mathematically correct, the program was inefficient in terms of both computational time and storage requirements. In its current implementation, the program is considerably more efficient and has an adequate front end which makes it accessible to a fairly wide class of users. Empirical results indicate that the size and complexity of problems that can be represented and solved with the system not only exceed the bounds of its original design, but are comparable and possibly superior to those of the best commercially available decision-analytic software.

Similarly, RACHEL's inference engine has been improved in several important ways. Prominent among these are a means for attaching general procedures at any point in the inference process, a variety of built-in procedures for the acquisition and display of information coupled with a facility for controlling these procedures (i.e., for the control of ASKability and TELLability), and a simple explanation mechanism.

C.2 Research in progress

The RACHEL system continues to be developed along four distinct directions: the efficiency and flexibility of RACHEL's inference engine are being improved, its explanation mechanism is being enhanced, RACHEL's facility for the development of patient preference models is being upgraded, and its knowledge base is being enlarged.

As it is currently implemented, the inference engine used by RACHEL is quite inefficient. This inefficiency is, to some extent, a deliberate design choice since the engine was designed to be very general and highly modular. Thus, there are many procedural redundancies and much unnecessary baggage in the programs that implement it. Now that we have a clearer idea of how the engine is to be used we have redesigned it by doing away with some of the original generality and modularity in favor of a more efficient process. Furthermore, the new design emphasizes and enhances particularly useful engine features such as its ASKability and its TELLability.

A further enhancement to RACHEL's inference engine concentrates on the system's ability to explain its line of reasoning. The original design only responds to online "why" queries by displaying its dynamic goal stack. In its new form, the engine allows offline as well as online queries in both "why" and "how" formats.

Beyond traditional explanation capabilities, we are exploring possible means to explain decision-theoretic inferences. In particular, we are trying to understand how to explain decision recommendations that are based on the maximization of expected utility to users unfamiliar with decision theory. Our current research indicates that a promising way to do this is to break down large decision problems into smaller, more manageable pieces whose formal solution can be checked against intuition. Although still at an early stage, this line of research seems to be on the path of eliminating an important barrier to the widespread use of normative decision techniques.

An exciting area of current interest is the improvement of RACHEL's facility for the creation and assessment of parametric models of patient preferences. In particular, we are trying to increase the generality of RACHEL's model library to account for acute as well as chronic conditions and to simplify the corresponding assessment process. This simplification is based on the notion that a better understanding of the major concerns of patients can help us redesign the questions asked by RACHEL so that they are closer to the specific experiences of individual patients. As part of this effort, we expect to have significant contact with actual patients to ensure the clinical relevance of our research.

A fourth area where RACHEL is being enhanced is the expansion of its medical and decision-analytic knowledge bases. Planned additions include further knowledge about the treatment of tubal blockage (including more data on in-vitro fertilization procedures and an ability to consider a wider class of patients) and a new packet of knowledge dealing with deterministic sensitivity analysis.

In addition to the development of RACHEL, there are several active research programs within the CAMDA project. One such program is aimed at developing a representation for dynamic decision processes (such as those faced by cancer patients) that do not necessarily satisfy the Markov assumption. This research has led to a generalization of influence diagrams which allows multiple value nodes. This generalization makes it possible for complex sequential decision processes (whose solution would otherwise be infeasible) to be efficiently solved.

Another research program within the CAMDA project is the development of fast algorithms for the solution of decision problems formulated as influence diagrams. In general, the solution of an influence diagram (i.e., the calculation of a recommended decision strategy) is obtained by the repeated application of an operation, known as "removal", to all nodes in the diagram other than the value node. The removal of a node in the diagram is a generalization of the foldback operation needed to solve a decision tree. With rare exceptions, the order in which nodes are removed from a diagram is not unique. Current results indicate that significant reductions in the computational burden of solution can be achieved by controlling the order in which diagram nodes are selected for removal.

At a more fundamental level, we are exploring the consolidation of the predicate calculus with probabilistic logic. Of particular interest is the design of an integrated inference engine that performs logical inferences within a probabilistic framework. A central problem in this research is the definition of universal and existential quantification in probabilistic terms.

A recent addition to our research project is a program to design cost-effective strategies for monitoring the recurrence of bladder cancer. We expect this research to interact with our ongoing search for more effective models of patient preferences.

D. Publications

1. Holtzman, S.:*A Model of the Decision Analysis Process*, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1981.

2. Holtzman, S.:*A Decision Aid for Patients with End-Stage Renal Disease*, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1983.

3. Holtzman, S.:*On the Use of Formal Models in Decision Making*, Proc. TIMS/ORSA Joint Nat. Mtg., San Francisco, May, 1984.

4. (*) Holtzman, S.: *Intelligent Decision Systems*, Ph.D. Dissertation, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1985.

5. Shachter, R.: *Evaluating Influence Diagrams*, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1984.

6. Shachter, R.: *Automating Probabilistic Inference*, Department of Engineering-Economic Systems, Stanford University, Stanford, California, 1984.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

II.A Medical Collaborations and Program Dissemination Via SUMEX

Since its inception, the CAMDA project has benefited from an active relationship among decision analysts, computer scientists, and members of the Stanford medical community. In particular, RACHEL is being developed in close cooperation with physicians in the Infertility Clinic at Stanford. Other programs within the CAMDA project such as our research on the form and use of medical preference models are being done in cooperation with physicians at the Palo Alto Veterans Administration Hospital and at El Camino Hospital.

II.B. Sharing and Interactions with other SUMEX-AIM Projects

II.B.1 SUMEX-AIM 1984 Workshop:

Samuel Holtzman participated in the 1984 AIM workshop in Columbus, Ohio. In addition to the presentation of a summary of CAMDA research, he had many opportunities to interact with workshop participants on an informal basis. Of particular interest were several discussions with members of the MIT/TUFTS group interested in medical decision analysis which have led to an interchange of ideas that continues to this date.

II.B.2 Decision Systems Laboratory Research Meetings

As part of the CAMDA project, we have instituted a weekly research meeting for those interested in the design and implementation of computer-based decision systems. This weekly meeting has become a very active forum for the presentation of research results. The following topics of direct relevance to medical decision making were presented during the last two academic quarters.

Date	Speaker	Topic
03-OCT-84	Ross Shachter	Probabilistic Inference
17-OCT-84	Jack Breese	Dempster-Shafer Theory
24-OCT-84	Kazuo Ezawa	Efficiency in Solving Influence Diagrams
07-NOV-84	Majid Khorram	Fuzzy Sets and Decision Making
14-NOV-84	Dan Kent	Utility Theory Underlying Physicians' Treatment Thresholds: HELP!
21-NOV-84	Yann Bonduelle	Explanation in Decision Systems
09-JAN-85	Ross Shachter	What Do You Call the Offspring of SUPERID and INFLUENCE?
23-JAN-85	Doug Logan	The Value of Probability Assessment
06-FEB-85	Seok Hui Ng	Minimal Tumor Follow-up Examination Schedule for Recurrent Bladder Cancer Patients.
13-FEB-85	Keh-Shiou Leu	TEREISIAS' Explanation Facility
06-MAR-85	Joe Tatman	Algorithm for Decision Processes Optimization
13-MAR-85	Gerald Liu (UC)	Knowledge Structure in Evidential Reasoning

II.B.3 Course in Medical Decision Analysis

A new course in medical decision analysis, taught by Prof. Samuel Holtzman, is being offered for the first time during the Spring quarter of 1985. The course is offered jointly by the Engineering-Economic Systems Department, the Medical Information

Sciences Program, and the Computer Science Department. The objective of the course is to expose students to the practice of decision analysis for clinical purposes and to introduce them to the design and use of computer-based medical decision tools.

II.C. Critique of Resource Management

The CAMDA project is heavily dependent upon the availability of the SUMEX computing resource. The physical facility as well as the staff of SUMEX-AIM are excellent. In particular, it has been a pleasure to deal with Ed Pattermann, who is invariably courteous, responsive to our needs, and effective in his actions. We will certainly miss him now that he has moved to industry. Pam Ryalls has also provided much needed help in managing the CAMDA project in a manner that is friendly and efficient.

As an update to last year's report, the previously reported Ethernet deficiencies have been corrected. This improvement was part of a campus-wide effort to improve Stanford's computer network which directly affected our campus connection to SUMEX. The system load on SUMEX continues to be heavy, although it appears to be somewhat lower than it was last year. The ability of the CAMDA project to use the DECSYSTEM-2020 machine operated by SUMEX (referred to as TINY) has had a significant effect on our ability to demonstrate our systems during normal business hours, further reducing our frustration with the main system's load.

III. RESEARCH PLANS

III.A Project Goals and Plans

During the upcoming year, we intend to enhance four specific elements of the RACHEL system: its inference mechanism, its explanation facility, its ability to model patient preferences, and its medical and decision-analytic knowledge bases. Furthermore, we intend to continue to improve our understanding of normative decision methodologies, with particular emphasis on the use of these methodologies for computer-based decision support. Section I.C.2 describes the near-term goals of the CAMDA project in more detail. Our long-term goal remains that of designing and implementing usable, fully-validated and documented systems for medical decision support.

III.B Justification and Requirements for Continued SUMEX Use

The CAMDA project is truly interdisciplinary. It draws on elements of decision analysis, artificial intelligence, and medical science. The project has the potential to contribute to each of these disciplines in important ways.

In particular, the CAMDA project is likely to lead to the development of tools and techniques that greatly improve the quality of decision making in medicine. For instance, RACHEL explicitly considers uncertainty, decision alternatives, and patient preferences in developing recommendations. In spite of its generality, RACHEL's interaction with the user is sufficiently terse and simple to support the claim that systems based on its methodology can be effective clinical decision tools. Much of the simplicity and terseness of RACHEL's operation is a direct consequence of the AI foundations of the system's design.

The heavy reliance of the CAMDA effort on artificial intelligence technology make SUMEX-AIM an ideal environment in which to pursue this research.

III.C Needs and Plans for other Computing Resources beyond SUMEX-AIM

The CAMDA project has access to four Olivetti M24 and one MAD-1 personal computers (IBM-PC type) as well as to one Apple Macintosh (128K) computer. In addition, we continue to search for funds to acquire one or more state-of-the-art LISP machines.

III.D Recommendations for Future Community and Resource Development

What would be the effect of imposing fees for using SUMEX resources (computing and communications) if NIH were to require this?

A major benefit provided by the existing SUMEX-AIM facility is the availability of very low-cost computing resources. Access to these resources is granted primarily on the basis of an assessment of the value of the proposed research to the overall goal of making artificial intelligence a useful medical tool. Imposing fees for using SUMEX would prevent users with modest means from obtaining access to the facility on the basis of merit alone.

Do you have plans to move your work to another machine workstation and if so, when and to what kind of system?

The CAMDA project has access to several personal computers for its research. These machines include Olivetti M24's (marketed as the A.T.&T. personal computer in the U.S.) and a MAD-1 personal computer -- all of which are compatible with the IBM-PC. In addition, the project has purchased an Apple Macintosh. These machines are used as a supplement to the SUMEX mainframe, and are not intended to replace it.

IV.C.2. Protein Secondary Structure Project

Protein Secondary Structure Project

Robert M. Abarbanel, M.D., Ph.D.
Section on Medical Information Science
University of California Medical Center
University of California at San Francisco

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

Development of a protein structure knowledge base and tools for manipulation of that knowledge to aid in the investigation of new structures. System to include cooperating knowledge sources that work under the guidance of other system drivers to find solutions to protein structure problems. Evaluations of structure predictions using known proteins and other user feedbacks available to aid user in developing new methods of prediction.

B. Medical Relevance and Collaboration

Many important proteins have been sequenced but have not, as yet, had their secondary or tertiary structures revealed. The systems developed here would aid medical scientists in the search for particular configurations, for example, around the active sites in enzymes. Predictions of secondary structure will aid in the determination of the full "natural" configuration of important biological materials. Development of systems such as these will contribute to our knowledge of medical scientific data representation and retrieval.

C. Highlights of Research Progress

The prediction of beta-alpha protein structures was completed in 1982. The system was developed on a VAX 11/750 at the University of California, San Francisco, to allow researchers to describe patterns of amino acid residues that will be sought in the sequences under study. The presence or absence of these "primary" patterns was then combined with other measures of structure, like hydrophobicity, to suggest possible alpha helix or beta sheet or turn configurations.

The segments of a sequence between turns were then analyzed to determine the allowable extent of the possible secondary structure assignments. Any segments remaining were then used to generate all possible complete structures. Only two beta strands with the character of sheet edges are allowed in any prediction. This hierarchical generation and pruning resulted in nearly 95% turn prediction accuracy, and excellent delimiting of helices and sheets. In some cases, one and only one secondary structure was predicted.

Work in progress -- The original pattern matching and manipulation system written in the C language, was re-written in Franz Lisp to run under UNIX(TM). This system was then re-written to run under KEE, The Knowledge Engineering Environment (TM, Intellicorp) on a Symbolics 3600 at the Computer Graphics Laboratory, University of California, San Francisco. This environment provides for ready display of pattern matches and viewing and manipulating of the applications of sets of rules. The original α/β rules are being tested now and new sets of rules are under development for